

Exhibit J

30 April 2010

Use and Misuse of Event Studies to Examine Market Efficiency



Dr. David Tabak

Senior Vice President
NERA Economic Consulting

Proof of market efficiency is often important for securities class actions at both the class certification and merits stages. Under the Supreme Court's ruling in *Basic v. Levinson*, if a market is efficient, there is a rebuttable presumption that investors relied on market prices, meaning that there is no need for each potential plaintiff to demonstrate individual reliance on any misrepresentation or omission, as those would be reflected in the market price. Courts have recommended various tests, or factors to consider, to assess whether a financial market is efficient. Unfortunately, the primary test that is designed to be consistent with the academic literature on market efficiency, an examination of whether prices respond to news, is often performed in ways that render it meaningless. This paper reviews how this test has been implemented and discusses what different versions of the test can reveal about market efficiency. It also discusses recent decisions that directly or indirectly recognize the potential pitfalls in improper ways of performing this test.

Market Efficiency in the Courts

In *Basic v. Levinson*,¹ the United States Supreme Court provided plaintiffs in securities class actions with a rebuttable assumption of reliance if the security that is the subject of the suit traded in an efficient market. Analysis of market efficiency has been performed at both the class certification and merits stages of securities fraud cases.²

The seminal case on testing for market efficiency, *Cammer v. Bloom*, identified five factors for which the court felt analysis would be helpful in alleging market efficiency: high weekly turnover, numerous market makers and arbitrageurs, the degree of analyst coverage, the issuer's eligibility to file an S-3 registration statement, and the showing of an empirical relationship between new information and movements in the price of a security.³ Analyses of these factors have often formed the backbone of tests of market efficiency in securities class actions, though they have since been expanded upon by cases such as *Krogman*.⁴

The *Cammer* Court noted that one factor, an empirical response of prices to new information was “the essence of an efficient market and the foundation for the fraud on the market theory.” Other courts have recognized the supremacy of this factor as well. See, for example, *PolyMedica*, in which the Court stated that the cause-and-effect relationship “is ‘in many ways, the most important,’ [citing *In re Xcelera.com Securities Litigation*]...”⁵ To a financial economist, this is indeed correct, as an efficient market is defined as one that quickly and fully incorporates all information of a certain type (pricing information for the security in weak-form efficiency; all public information for semi-strong-form efficiency; and all information for strong-form efficiency).⁶

The other factors cited by courts are designed to be conditions that are likely to either be conducive to or the result of an efficient market. For example, because analysts provide information and insight to market participants, their presence might be considered helpful to the functioning of an efficient market. A contrary view is taken by some academics who note that many of those other factors might appear because a market is inefficient.⁷ For example, there may be a greater need for, and thus a greater presence of, analysts in a market that is not efficient rather than one in which market participants can rely on market prices. Also notable is that many of these factors have no bright line dividing results that favor efficiency versus those that favor inefficiency.

Empirical efforts to examine whether other factors are associated with market efficiency have led to mixed results. For example, one early examination of these factors found that trading volume and the number of analysts is correlated with market efficiency, while other factors, including market capitalization, number of market makers, bid-ask spread, and institutional holdings, are not.⁸ Notably, this study categorizes stocks as efficient or not by the movement of the stock price in response to earnings information, or an empirical response of prices to news. Later studies similarly find that many of the factors used by courts do not appear to correlate with the results of empirical measurements of market efficiency. Thus, while these factors are often considered important by courts, there is little academic support or analysis of how they should be performed. Consequently, this paper focuses on the analysis of price responses to news, while recognizing that some courts may consider this factor among others in their analysis of market efficiency.

How Has the Price Response to News Been Examined?

There are various ways in which one can attempt to examine whether stock prices respond to news. Perhaps the easiest, and the least reliable, is a form of “proof by example,” in which one finds a number of days in which there is news followed by a relatively large stock price movement. The court in *PolyMedica*, properly noted that the “mere listing of five days on which news was released and which exhibited large price fluctuations proves nothing.” Though this proof-by-example may be fancied up with various statistics (e.g., making sure that the stock price movement is statistically significant under some form of measurement) or even by looking at all of the days with price movements at or above some level, this form of proof by example still proves nothing, or perhaps more accurately proves merely that there were a few days when there was news and a large price movement without proving any reliable correlation between those two factors. Also proving nothing are other *non-comparative* analyses such as determining the percent of all days with a significant price change that were associated with news or the percent of all days with news that were associated with a significant price change.

To see why these analyses are not meaningful, consider an example based on the table below:

Day	News	Significant Price Movement
1	Yes	Yes
2	Yes	Yes
3	Yes	Yes
4	Yes	Yes
5	Yes	No
6	Yes	No
7	Yes	No
8	Yes	No
9	Yes	No
10	Yes	No
11	Yes	No
12	Yes	No
13	Yes	No
14	Yes	No
15	Yes	No
16	Yes	No
17	No	Yes
18	No	No
19	No	No
20	No	No

In this example, there are twenty days, divided as follows:

- Four days (Days 1-4) have news and a significant price movement
- Twelve days (Days 5-16) have news, but not a significant price movement
- One day (Day 17) has no news, but has a significant price movement
- Three days (Days 18-20) have no news and also do not have a significant price movement.

Notably, only four of the sixteen days with news, or 25%, are associated with a significant price movement. Of the four days without news, only one, or 25%, is associated with a significant price movement. Because the same percent of the days with and without news is associated with a significant price movement, one cannot say that news is correlated with a significant price movement.

Under an improper form of analysis, one would begin by identifying the five days with significant price movements (Days 1-4 and 17). At the most basic level, one might just mention Days 1 through 4 and claim that this listing of four days with price movements and news proves some form of association. Being a little more systematic, the proponent of market efficiency could argue that four of the five days, or 80%, with substantial price movements are associated with news. He would then incorrectly conclude that significant price movements are correlated with news.

If the 80% figure were used to suggest that news is highly likely to cause a price movement, this would be an example of the Prosecutor's Fallacy.⁹ Having determined that significant price movements have an 80% likelihood of being associated with news, one may implicitly suggest that news has a high (such as 80%) likelihood of generating significant price movements. One way to see this fallacy is to note that on 100% of the large price movement days, the sun did indeed rise. This would not, however, mean that the sun's rising causes large price movements. Failing to consider all of the data, but instead focusing on only those days with significant price movements, will have led to a meaningless analysis.

To see this from a statistical sense, consider what the error rate for such an analysis represents. Answering this question requires us to first determine the end point of the analysis. Suppose that one's goal is to find at least one day with a significant price movement and news. If price movements are completely uncorrelated with news, the probability of not finding such a day in a random sample of N days is $[1 - \alpha * p]^N$, where α is the proportion of days with significant price movements (e.g., 0.05), p is the probability that there is news on any day (e.g., 0.40). If N is equal to 10 days, then the probability of finding at least one example of a day with both news and a significant price movement (i.e., the false-positive error rate) is 0.18, which is not the same as α , the share of price movements that is deemed to be significant. (If an event study is used to determine significance, α would be the cutoff level for statistical significance; hence the α of 0.05 corresponds to the standard 5% significance level, and similarly corresponds to the probability of a false positive.) Consequently, a proof-by-example methodology has a much different, and typically much higher, error rate than does a standard event study.

How Should One Measure the Responsiveness of Prices to News?

Given that the proof-by-example types of methodologies do not provide reliable evidence in favor of market efficiency, we turn to more appropriate approaches. Courts have commonly looked to an "event study" to determine whether prices respond to corrective disclosures and should look favorably on similar analyses of whether there is a systematic relationship between news and price movements.¹⁰ Using an event study to make this determination is not just in accord with prior courts' decisions, but fits in with a statement by Eugene Fama, often considered the father of the efficient market hypothesis, that "event studies give the most direct evidence on efficiency."¹¹ Indeed, the academic literature is filled with many properly performed event studies that shed light on how stocks do or do not respond to certain types of news. The degree to which those responses are consistent with theory (e.g., price responses to new material news but not to repetitions of old news) provides evidence relevant to an examination of market efficiency.

One interesting question that has arisen in various cases is the proper order of steps in an event study. In particular, some experts have stated that one must first find the news and then measure the price movements while others have argued that one can first identify the significant price movements and then assess the news.

Fortunately, the relevant academic literature is clear: first one identifies news and then measures price movements associated with that news.¹² The reasoning for this is rooted in the basics of scientific methodology: first a hypothesis is proposed and then a test is made to see if that hypothesis can be rejected. The statistics of an event study apply only if the statistical test is made *after* the hypothesis is created, so that there is a chance that the price movement is statistically significant and a chance that it is not. If one first identifies days that have statistically significant movements, then the only error can occur in properly identifying material news. However, this is not the function of an event study, and certainly the error rates from the event study do not apply to the probability of making accurate news determinations.¹³

Thus, any process that first looks at price movements and then searches for news is not a proper event study, at least not if news determinations involve any element of subjectivity. But, could it still be a proper way to examine market efficiency? Certainly the subjective determination of what is news after knowing the magnitude of the associated price movements is counter to sound scientific practice. Unfortunately, one court recently rejected defendants' claim that plaintiff's event study was wrong because, among other things, "the study first identified statistically relevant price movements and then thereafter identified [sic] corresponding news events."¹⁴ Defendants' objection was, in fact, completely correct. The court, however, dismissed the objection, arguing that "a study that first focused on news events and only then attempted to analyze price fluctuations would be ambiguous: if a stock price was seemingly unaffected on the date of a news release, one would not be able to discern whether this was due to market inefficiency or simply investor indifference to that particular news event." While true, any such ambiguity would occur whether the news was identified before or after the price response was measured. If anything, making such a subjective analysis after knowing the results of the statistical examination of the price movement is more likely to result in an improper and potentially biased analysis.¹⁵

Another problem with identifying news after identifying meaningful price movements is that it is difficult to say what would have happened if additional price movements were analyzed, for example by having a laxer standard for what counts as a relevant price movement or under a different model for adjusting for market or industry effects. Only on occasion is this possible, such as in the recent *In re AIG Securities Litigation* opinion, in which the Court was able to examine evidence that showed that "on one of the key dates during which significant AIG-related disclosures occurred, there was little or no movement in the price of the debt securities, despite large decreases in the price of AIG stock, both in absolute terms, and in the amounts attributable to the disclosures according to [the expert's] event study."¹⁶ It was only because the expert had performed event studies on the company's stock, and in so doing opined that news on certain days was relevant, that there was a set of news days that the expert had accepted as material for which movements of the debt securities could be examined. Notably, because there was no movement on the date mentioned above (as well as no trading, and thus no price movement, on other days, as discussed in the opinion), one could see how the debt securities failed to respond to news that the expert felt was material; this is generally not possible if one identifies news on only dates considered to be material.

Furthermore, though this may be a point of contention in an actual litigation, suppose for purposes of illustration that one can clearly distinguish material (meaning important and containing new information) news from non-material news in an objective fashion. Doing so after examining the price movements provides no benefit to a properly conducted examination of market efficiency because all news events in the study, whether associated with significant or insignificant price movements, must be characterized in such an analysis. Doing so before examining the statistical results can aid in minimizing debates about whether the nature of the manner in which material news was identified and whether that method is indeed truly objective.

Assuming that the news has been properly categorized and the price movements measured, the analysis can, in fact, be treated as analogous to a form of epidemiological study. Epidemiology generally studies how environmental, genetic, and other factors are associated with, and may potentially cause, some disease. Here, the potentially causative factors under investigation are news events and the outcomes of interest are large price movements.¹⁷

The *Reference Guide on Epidemiology*, a part of the *Reference Manual on Scientific Evidence*, published by the Federal Judicial Center, provides a useful discussion of different types of epidemiological studies. Relevant for our purposes is the discussion on pages 340 and 341, which notes that to determine whether exposure to a potential causal factor (in our case, exposure of the market to new information) is related to an outcome (there, a health outcome; here, a significant price change), one can create a cross-tabulation of exposure as follows:

Cross-Tabulation of Exposure by Disease Status

	No Disease	Disease
Not Exposed	a	b
Exposed	c	d

According to the *Reference Guide on Epidemiology*, using the data in such a table, the “researcher would compare the portion of unexposed individuals (controls) with the disease ($b/(a + b)$) with the proportion of exposed individuals (cohort) with the disease ($d/(c + d)$). If the exposure causes the disease, the researcher would expect a greater proportion of the exposed individuals than of the unexposed individuals to develop the disease.”¹⁸

Hence, analyses that use all four variables in Table 1 from the *Reference Guide on Epidemiology* may, if implemented properly, shed light on the question of whether news is associated with significant price movements. Consider the following table, the same as in the *Reference Guide on Epidemiology*, but with the names of the rows and columns altered.

	No Change In Price	Change In Price
No News	a	b
News	c	d

The improper analyses examined above do not use all four variables and do not compare ($b/(a + b)$) with ($d/(c + d)$). That is, they do not compare the fraction of days with no news that exhibit a significant price change to the fraction of days with news that exhibit a significant price change. The proof-by-example methodology is to show that some days that are “exposed” to news have a resulting development (i.e., a significant price movement, the analogue to the disease). That is, it simply provides a few examples to show that d is not zero. This is true whether the elements of d are just generally large or are days with statistically significant price movements. An examination of all the days that have large price movements (i.e., those days in b and d), even if it shows that d is large relative to b+d (i.e., that a large fraction of the total days with significant price movements is associated with news) does not perform the relevant test indicated by the *Reference Guide*, and fails to present any scientifically relevant evidence.

The first step in a proper test

Faced with an offer of proof of a price response to news via five selected examples, the court in *PolyMedica* said that to “approach usefulness, an analysis should statistically compare all news days with all non-news days.” While the *PolyMedica* court was correct that a systematic examination of days is needed, the burden it imposed of looking at all days can be lessened if an analyst can create a valid sample of days that is selected in a way that is not based on

price movements. Still, the general principle behind the reasoning of the *PolyMedica* court, that one must engage in a systematic study and not solely select examples that go in a certain direction, is important.¹⁹

An example of how to perform such a test is provided in Ferillo, Dunbar, and Tabak (FDT) in *St. John's Law Review*. The authors describe a methodology in which news days and non-news days are categorized into those that are associated with statistically significant returns and those with statistically insignificant returns. The proportions of each type of day (news or non-news) with statistically significant returns are compared to see if they differ by a statistically significant amount. If the news days are more likely to be associated with a statistically significant return, this would be evidence in favor of market efficiency, while if the difference in proportions is not statistically significant (or if the non-news days are more associated with statistically significant returns), this is a statistical rejection of market efficiency.

Note that the FDT methodology can be altered to compare the average absolute value of the market-adjusted returns for the news days and the non-news days, and not just the percent of days that have statistically significant returns. In addition, one could implement another version of the methodology by measuring the excess returns and including a dummy variable on each of the news days. If the dummy variables are jointly significant, that would be evidence that, on average, there was some response to news.

Where this method falls short

There are several ways that versions of the FDT methodology may not be able to fully distinguish an efficient market from an inefficient one.

To begin, there are questions of implementation: e.g., To what extent has news been properly identified, and what happens if the news is not properly identified? If a substantial fraction of non-news days are misclassified as news days (e.g., if repeated news or immaterial news events are counted as news), the average price movements across news and non-news days should move closer together, making it harder to distinguish them. Similar results apply if news days are misclassified as non-news days. If both errors occur, the question becomes whether the relative compositions of the two groups (purported news days and purported non-news days) are still different enough for the test to be valid. (E.g., at an extreme where both contain the same percentage of true news and true non-news days, one would not expect to see any difference in price behavior even if the market is perfectly efficient.) These types of problems presumably would plague all tests of price response to news, since one cannot perform such a test reliably without knowing what is and what is not news.

There are also conceptual questions relating to whether a market can exhibit some form of inefficiency but still pass the FDT test. In fact, there are clear examples where that could be the case.

First, what if only some news items are incorporated into the price of a security? The result could look like efficiency but the security may not fully impound *all* news. For example, suppose that on non-news days a stock moves by 2% in absolute value and on news days it moves by 8%. If the stock fails to respond to half of the news, we would observe average movements of 2% in absolute value in the non-news sample and of 5% in the news sample (the 5% being the average of the half of the news days when the stock moves by 8% because it responds to news and the half of the news days when it moves by the 2% typical of the movement on a non-news day). The price movements in the two samples look different. This does show that the

stock sometimes does respond to news, but not that there was a “reliable relationship between changes in [the company’s] stock price and news events.”²⁰ Consequently, we cannot be sure that the market would incorporate any misrepresentation, meaning that investors should not rely on market prices.

Second, what if news is not fully impounded immediately? Then excess price movements would be likely to occur on days following news days (which themselves could be random). While the stock would pass the initial test of whether price movements on news days and non-news days look different, it may fail a different type of test, such as whether there is any persistence or serial correlation in returns. Note that the question here is not solely whether the stock takes some time to fully impound news (which is inconsistent with a strict reading of an efficient market’s responding quickly to news), but also whether the time pattern in which news is impounded is predictable (which is inconsistent with any view of an efficient market).

Different forms of the FDT test are more susceptible to different types of inefficiency. For example, suppose that a stock was generally unresponsive to news but did respond to the most extreme news events, such as news that the company was nearly worthless. In that case, the stock movement on most of the news days would be indistinguishable from the stock movement on non-news days, which would mean that the stock was *inefficient*, because it did not incorporate all news. Given a large enough sample and only infrequent occurrences of such major news, a comparison of the percentage of news and non-news days with statistically significant returns should correctly provide evidence against market efficiency. However, given the large price movement on those few days with highly material news, tests of the average response to news or a test of the price movements on news days as a group may find that there is some response to news. In effect, the tests are addressing different questions: the comparison of the proportion of news days and non-news days examines whether the stock consistently responds to news, while the other tests address the question of whether the stock responds on average to news, even if that average response is driven by a small subset of the news events. If the test is designed to examine whether all news is fully impounded into a stock price, then in this instance the proportions test gives the correct answer, while others do not.

An example in the other direction may be when the typical news content is much less than the volatility in the stock price. In that case, the distribution of returns on news days is likely to look similar to the distribution on non-news days, but have a slightly larger variance. If that increase in variance is proportionately small, few news days will make it across a division between statistically insignificant and statistically significant returns. On the other hand, an examination of the average absolute return or an F-test on the returns of news days is more likely to pick up this effect, thereby finding that there is some response to news, though without showing that these are reasonably sized responses to the news events.

Conclusion

There are various ways to test if the price of a security responds to news. For any such test to yield valid evidence in favor of market efficiency, however, it cannot be some form of proof by example or an examination of only days with large price movements. A showing of market efficiency will generally require both a treatment group and a control group, with the treatment group being days with news and the control group days without news. Other forms of analysis may be possible, but should generally involve some form of comparison between two groups with a statistical test to see if that difference is meaningful.

Consistent with the scientific method, one never proves market efficiency. Instead, one begins with efficiency as a hypothesis and sees if there is sufficient evidence to reject that hypothesis. One potential concern is that this gives those opposed to market efficiency multiple bites at the apple should they decide to keep testing until they find one test that rejects market efficiency. Here, both statistics and common sense may be helpful. Statistically, one can and should recognize that as more tests are performed, the likelihood of a false positive goes up. There are sometimes ways to adjust for this and, when appropriate, they should be applied.²¹ Common sense also plays a role. It is one thing to argue that a stock does not react efficiently because a portion of the price movement on one day is, on average, continued or reversed on the following trading day. It is another to argue that a portion of the price movement on one day is, again on average, continued or reversed eight trading days later with no statistically significant relationship between the initial movement and the seven interim ones. A response on the day after news sounds like a potential effect of an inefficient market that incorporates news slowly; a response only eight trading days later sounds like the result of a search over various possible lag structures.

While we may never be able to definitively show market efficiency or inefficiency, there are stronger and weaker tests, and even tests that prove nothing. By considering the reasonableness of each test, particularly those that have easy counter-examples such as what was shown above, experts, parties, and courts are more likely to reach conclusions based on those tests that are in fact reliable.

End Notes

- 1 *Basic v. Levinson*, 485 U.S. 224, 108 S.Ct. 978 L.Ed.2d. 194 (1988).
- 2 The Ninth Circuit has recently held that "Plaintiffs pleading fraud-on-the-market, on the other hand, may have to establish an efficient market to even raise common questions or show predominance." (*Dukes v. Wal-Mart Stores, Inc.*, Nos. 04-16688 & -16720 (9th Cir. Apr. 26, 2010))
- 3 *Cammer v. Bloom*, 711 F. Supp 1264 (D.N.J. 1989).
- 4 *Krogman v. Sterritt*, 202 F.R.D. 467 (N.D. Tex. 2001).
- 5 *In re PolyMedica Corp. Sec. Lit.*, 453 F. Supp. 2d 260, 266 (D. Mass. 2006).
- 6 See, for example, the *New Palgrave Dictionary of Economics Online*, http://www.dictionaryofeconomics.com/article?id=pde2008_S000529&goto=stock&result_number=1671.
- 7 Grigori Erenburg, Janet Kiholm Smith, and Richard L. Smith, "The Paradox of 'Fraud-on-the-Market Theory': Who Relies on the Efficiency of Market Prices?", Working Paper, 12 November 2008, p. 3.
- 8 Brad M. Barber, Paul A. Griffin, and Baruch Lev, "The Fraud-on-the-Market Theory and Indicators of Common Stocks' Efficiency," *The Journal of Corporation Law*, Winter 1994.
- 9 See, for example, C.G.G. Aitken, "Interpretation of Evidence, and Sample Size Determination," in Joseph L. Gastwirth, ed., *Statistical Science in the Courtroom*, 2000, p. 5.
- 10 While academic use considers an event study to be the statistical examination of an event or a type of event (e.g., announcements of stock splits), lawyers and courts in securities class actions have often used the term "event study" to mean either a chronology with news, prices, and calculations such as excess returns and t-statistics associated with those returns, or else for an in-depth analysis of all the news on a particular issuer on a particular day, ideally including the parsing of the effects of the different components of that news on the price of a security of the issuer.
- 11 Eugene Fama, "Efficient Capital Markets: II," *The Journal of Finance*, Vol. XLVI, No. 5, December 1991, p. 1577.
- 12 See, for example, Robert G. Bowman, "Understanding and Conducting Event Studies," *Journal of Business Finance & Accounting*, 1983. Bowman identifies four types of event studies, including "Market Efficiency" event studies. Bowman provides a list of steps for this type of event study, in which the first is "Identify the event of interest." He states, "The first step in the conduct of an event study is to identify the event of interest." (p. 563) See also Craig A. MacKinlay, "Event Studies in Economics and Finance," *Journal of Economic Literature*, Vol. XXXV (March 1997). ("The initial task of conducting an event study is to define the event of interest and identify the period over which the security prices of the firms involved in this event will be examined—the event window.") (p. 14)
- 13 Similarly, good medical research has the doctors blinded (i.e., not knowing whether they are evaluating patients that received a treatment or a placebo) when evaluating individual patients.
- 14 *In re DVI Inc. Securities Litigation*, 249 F.R.D. 196 (E.D. Pa. 2008).
- 15 A criticism of selecting dates subjectively is found in *Bell v. Ascendant Solutions, Inc.*, No. 01-0166, 2004 WL 1490009, at 3 (N.D. Tex. July 1, 2004), in which the court, in reviewing an expert's "information dates," found that he "include[d] dates that appear to be consciously chosen in order artificially to support his hypothesis of efficiency."
- 16 *In Re American International Group Securities Litigation*, No. 04 Civ. 8141 (DAB), 2010 U.S. Dist. LEXIS 15453, at *82-*105 (S.D.N.Y. Feb. 22, 2010). For purposes of disclosure, the author of this paper was a defense expert on this case until his clients settled before the cited opinion was issued.
- 17 Because news events vary in content, this is analogous not to an epidemiological study that covers very similar types of exposures, but to one that examines exposure to different, but presumably somewhat similar, causal factors. (E.g., a study on whether workers in chemical labs, with exposure to various chemicals, develop a certain disease, as opposed to a study of exposure to one particular chemical.)
- 18 *Reference Guide on Epidemiology*, p. 341, internal footnote omitted.
- 19 One counterexample is sufficient to disprove a hypothesis that something never happens. For example, a single coin flip of tails will disprove the hypothesis that that coin only comes up heads. Because an efficient market absorbs all news, one or else a limited number of counterexamples will disprove market efficiency (or at least require the period in which the market may have been efficient to be limited to a time frame that excludes those examples), if it can be shown that the relevant news was in fact new and sufficiently material to have affected the price in an efficient market. In contrast, a limited number of examples cannot prove market efficiency.
- 20 *Krogman* at 477.
- 21 The concept here, known as "multiple comparisons," is discussed in many statistical texts.

About NERA

NERA Economic Consulting (www.nera.com) is a global firm of experts dedicated to applying economic, finance, and quantitative principles to complex business and legal challenges. For half a century, NERA's economists have been creating strategies, studies, reports, expert testimony, and policy recommendations for government authorities and the world's leading law firms and corporations. We bring academic rigor, objectivity, and real world industry experience to bear on issues arising from competition, regulation, public policy, strategy, finance, and litigation.

NERA's clients value our ability to apply and communicate state-of-the-art approaches clearly and convincingly, our commitment to deliver unbiased findings, and our reputation for quality and independence. Our clients rely on the integrity and skills of our unparalleled team of economists and other experts backed by the resources and reliability of one of the world's largest economic consultancies. With its main office in New York City, NERA serves clients from more than 25 offices across North America, Europe, and Asia Pacific.

Contact

For further information and questions, please contact the author:

Dr. David Tabak

Senior Vice President
+1 212 345 2176
david.tabak@nera.com

The opinions expressed herein do not necessarily represent the views of NERA Economic Consulting or any other NERA consultant. Please do not cite without explicit permission from the author.